

MAUS: The MICE Analysis and User Software

MAUS Developers

E-mail: durga@fnal.gov

ABSTRACT: The Muon Ionization Cooling Experiment (MICE) has developed the MICE Analysis User Software (MAUS) to simulate and analyse experimental data. It serves as the primary codebase for the experiment, providing for offline batch simulation and reconstruction as well as online data quality checks. The software provides both traditional particle physics functionalities such as track reconstruction and particle identification, and accelerator physics functions such as calculating transfer matrices and emittances. The code is structured in a Map-Reduce framework to allow parallelization whether on a personal computer or during data-taking operations. MAUS allows users to develop in either Python or C++ and provides APIs for both. Various software engineering practices from industry are also used to ensure correct and maintainable physics code, which include unit, functional and integration tests, continuous integration and load testing, code reviews, and distributed version control systems. The software framework and the simulation and reconstruction capabilities are described.

KEYWORDS: MICE; Ionization Cooling; Software.

Contents

1. Introduction	1
1.1 The MICE Experiment	1
1.2 Software Requirements	1
2. MAUS	2
2.1 Architecture	2
2.2 Data structure	4
2.3 Testing	5
2.4 Data Flow	5
3. Monte Carlo	5
3.1 Beam generation	5
3.2 GEANT4	6
3.3 Geometry	6
3.4 Tracking, Field Maps and Beam Optics	7
3.5 Detector response and digitization	7
4. Reconstruction	7
4.1 Time of flight	8
4.2 Scintillating fiber trackers	8
4.3 KL calorimeter	9
4.4 Electron-muon ranger	9
4.5 Cherenkov	9
4.6 Global reconstruction	9
4.6.1 Global Track Matching	10
4.6.2 Global PID	10
4.7 Online reconstruction	13
5. Summary	13

1. Introduction

1.1 The MICE Experiment

The Muon Ionization Cooling Experiment (MICE) sited at the Rutherford Appleton Laboratory (RAL) will be the first demonstration of muon ionization cooling – the reduction of the phase-space of muon beams. Muon beam cooling is essential for future facilities based such as the Neutrino Factory of Muon Collider [2, 3]. The experiment is designed to be built and operated in a staged manner. In the first stage (Step I), the muon beamline was commissioned [?] and characterized [?]. The present step – Step IV – will study the change in normalized transverse emittance using lithium hydride and liquid hydrogen absorbers under various optical configurations.




Figure 1. Rendering of the next MICE Step IV configuration

1.2 Software Requirements

The MICE software must serve both accelerator physics and particle physics needs of the experiment. Traditional particle physics functionality includes reconstructing tracks, identifying particles, and simulating the response from various detectors, while the accelerator physics aspect includes computing transfer matrices and Twiss parameters and propagating beam envelopes. All of these require knowledge of the beamline, geometries of the detectors, knowledge of the magnetic fields, and functionality to reconstruct or simulate various detectors thus necessitating a single software scope. Given the complexity and the inherent time-scales of experiments, the need also arises to ensure long-term correctness and maintainability of the software used for the experiment.

??? TODO: expand?

2. MAUS

The MICE Analysis User Software (MAUS) [7] is the experiment's simulation, reconstruction, and analysis software framework and aims to provide capabilities to 1) perform a Monte Carlo (MC) simulation of the experiment, 2) reconstruct tracks and identify particles from simulations and real data, and 3) provide monitoring and diagnostics while running the experiment.

Installation is by a set of shell scripts with SCons as the build tool. The codebase is maintained with the GNU Bazaar revision control system and is hosted on Launchpad. MAUS has a number of dependencies on standard packages such as Python, ROOT [13] and GEANT4 which are built as "third party" external libraries during the installation process. The officially supported platform is Scientific Linux 6 though developers successfully build on Ubuntu, CentOS and OpenSUSE Linux distributions.

There are 5 different detectors in MAUS: 1) time-of-flight (TOF) scintillators, 2) threshold Cherenkov (Ckov) counters, 3) scintillating fiber trackers, 4) "KL" sampling calorimeter, and 5) an electron-muon ranger (EMR). Their simulation and reconstruction software are described in Section 4.

??? TODO: say more about detectors and their capabilities/purpose ??

2.1 Architecture

The architecture is inspired by the Map-Reduce [?] data flow in order to simplify the interfaces that developers have to follow and aid running the code in parallel. It was felt that Map-Reduce paral-

lelizes particle physics problems in a useful fashion but the API was simplified to have *transformers* and *mergers* instead of maps and reduces.

MAUS is written in a mixture of Python and C++. C++ is used for complex or low level algorithms where processing time is important while python is used for simple or high level algorithms where development time is a more stringent requirement. Developers are allowed to write in either Python or C++ and python bindings to C++ are handled through internal abstractions or SWIG.

MAUS has an Application Programmer Interface (API) that provides a framework on which developers can hang individual routines. The MAUS API framework provides MAUS developers with a well-defined environment for developing reconstruction code, while allowing independent development of the backend and code-sharing of common elements like error handling and data mangling. A *Module* is the basic building block of the MAUS API framework. Four types of module exist within MAUS:

1. **Inputters** generate input data either by reading data from files or sockets, or by generating an input beam.
2. **Mappers** modify the input data, for example by reconstructing signals from detectors, or tracking particles to generate MC hits
3. **Reducers** collate the mapped data and allow functionality that requires access to the entire data set
4. **Outputters** save the data either by streaming over a socket or writing data to disk.

There are some objects that sit outside the scope of this modular framework but are nevertheless required by several of the modules. For instance, knowledge of the detector geometries, magnetic fields, and calibrations are required by the reconstruction and simulation modules, and objects like electronics cabling maps are required to unpack data from the data acquisition (DAQ) source, and error handling functionality is required by all of the modules. All these objects are accessed through a static singleton *globals* class.

The principal event type is the *spill*. A single spill corresponds to data from the particle burst associated with a dip of the MICE target. A spill typically lasts *3ms* and contains several DAQ triggers. Data from a given trigger correspond to a single MICE *event*. In the language of the Input-Map-Reduce-Output framework above, an Input module creates an instance of spill data, a Map module processes the spill (reconstructing, simulating, etc), a Reduce module acts on a collection of spills when all the mappers finish, and finally an Output module stores the MAUS data structure to the output.

MAUS has two execution concepts. A *job* refers to a single execution of the code, while a *run* refers to the processing of data for a DAQ run or MC run. Since data are typically accessed from a single source and written to a single destination, Inputters and Outputters are initialized and destroyed at the beginning and end of a job. On the other hand, Mappers and Reducers are initialized at the beginning of a run in order to allow loading run-specific information such as electronic cabling maps, fields, and calibrations.

Developers are allowed to write modules in either Python or C++ and python bindings to C++ are handled through internal abstractions.. C++ is used for complex or low level algorithms where




Figure 2. A Map-Reduce-Merge framework

computation time is important while python is used for simple or high level algorithms where development time is a more important criterion. In practice, all the reconstruction modules are written in C++ but support is provided for legacy modules written in Python.

MAUS has an Application Programmer Interface (API) that provides developers with a well-defined environment for developing reconstruction code, while allowing independent development of the backend and code-sharing of common elements like error handling and data manipulation.

Data can be represented in two formats. The default data format is a ROOT binary and the secondary format is JSON[12]. This is an ascii data-tree format readable with any text editor. Specific JSON parsers are also available - for example, the python *json* module is available and comes prepackaged with MAUS. There is a third special data type that MAUS handles, the Image type used by Reducers to output images of monitoring histograms, efficiency plots, etc and this is available only in JSON format.

Modules can exchange data either as C++ or JSON object types. In python, the data's format can be changed by using a converter module and in C++, mappers are templated to a MAUS data type and the API then handles any necessary conversion to that type. During production deployment of the software it was found that there was a significant hit in the performance speed due to inherent slowness in (de)serializing JSON objects. Hence it was decided that modules for official reconstruction and simulations must exchange C++ objects as a default in order to minimize need for conversion between data types. However, data can still be output in JSON format and

developers find it extremely useful during debugging.

2.2 Data structure

The principal part of the MAUS data structure is a ROOT Tree each entry of which corresponds to the data associated with a spill. A spill can have three event types associated with it: a MC event *MCEvent* contains an array of data each member of which represents the MC of a single particle traversing the experiment, a reconstructed event *ReconEvent* contains an array of data each member of which corresponds to a particle event corresponding to a trigger, and a raw data event *DAQEvent* corresponds to the raw data readout. There are 5 different detectors in MAUS: 1) time-of-flight (TOF) scintillators, 2) threshold Cherenkov (Ckov) counters, 3) scintillating fiber trackers, 4) "KL" sampling calorimeter, and 5) an electron-muon ranger (EMR). Each of these detectors have several layers of reconstruction which can be broken into individual reconstruction modules.

The *MCEvent* is subdivided into sensitive detector hits (energy deposited, position, momentum) and information about the primary particle that created the hits in the various detectors. The *ReconEvent* and *DAQEvents* are subdivided by detector. ReconEvents contain reconstructed particle data for each detector and the trigger. There is an additional branch that contains global reconstruction output, that is the track fitting between detectors.

The data can be written in two formats. The main data format is a ROOT binary format. This requires the ROOT package to read and write, which is a standard analysis/plotting package in High Energy Physics and is installed by the MAUS build script. The secondary data format is JSON. This is an ascii data-tree format that in principle can be read by any text editor. Specific JSON parsers are also available - for example, the python *json* module is available and comes prepackaged with MAUS.

There is a final data type that MAUS handles, the Image type used by Reducers to output monitoring histograms, efficiency plots, etc and available only in JSON format.

2.3 Testing

MAUS has a set of tests at the unit level and integration level. Unit tests are implemented against a single function, while integration tests operate against a complete workflow. Unit tests check that each function operates as intended by the developer and can achieve a high level of coverage and good test complexity. Integration tests check that the overall design of the code meets the specifications laid out, that interfaces with external codes or systems operate correctly.

The MAUS team aims to provide unit test coverage that executes 70–80 % of the total code base. This level of test coverage typically results in a code that performs the major workflows without any problem, but has errors in some of the less well-used functions and can behave ungracefully following user error. At the most recent release, MAUS test coverage was 68 % for python code, 78 % for non-legacy C++ code and 36 % for legacy (pre-2011) C++ code.

MAUS operates a continuous integration stack using a pair of test servers that mimic an online (control room) and an offline environment. Build and test is driven by the Jenkins test environment [8]. Developers are asked to perform a build and test on a personal code branch, using the test server, before integrating with the development trunk. This enables MAUS to make frequent clean releases. Typically MAUS works on a 2-4 week release cycle.

2.4 Data Flow

??? TODO: Need better (cleaner) data flow pictogram

The data flow of the MC and reconstruction algorithms is shown in Fig. 4. The data flow is grouped into three principal areas; the MC data flow is used to generate digits (electronics signals) from tracking; the reconstruction data flow is used to generate digits from raw data and the digits are converted to physics parameters of interest by the reconstruction.

3. Monte Carlo

An MC simulation of MICE encompasses beam generation, geometrical description of detectors and fields, tracking of particles through detectors, and digitization of the detectors' response to particle interactions.

3.1 Beam generation

The simulation within MAUS starts from the upstream end of the D2 magnet. The composition and phase space description of beam particles can be specified with datacards or can be read from an external file in The G4Beamline package is installed by the MAUS installation scripts and allows This To generate an incident beam, routines are provided to sample particles from a multivariate gaussian distribution or generate ensembles of identical particles ("pencil" beams). Additionally it is possible to produce time distributions that are either rectangular or triangular in time to give a simplistic representation of the MICE time distribution. Parameters, controlled by datacards, are available to control random seed generation, relative weighting of particle species, and the transverse-longitudinal coupling in the beam. Beam particles can also be read in from an external file G4Beamline [9], ICOOL [10], and user-defined formats. MAUS also allows the generation of a polarized beam by generating a spin vector from beam distributions.

In order to generate beams which are more realistic taking into account the geometry and fields of the actual MICE beamline, scripts to install G4Beamline are shipped with MAUS. This allows users to generate beam particles using G4Beamline and later simulate their tracking and interactions using MAUS.

??? TODO: more about g4bl generation??

3.2 GEANT4

The GEANT4 bindings are encoded in the Simulation module. GEANT4 groups particles by run, event and track. A GEANT4 run maps to a MICE spill; a GEANT4 event maps to a single inbound particle from the beamline; and a GEANT4 track corresponds to a single particle in the experiment.

GEANT provides a variety of reference physics lists to model the interactions of particles with matter. The default list in MAUS is *QGSP_BERT* which models hadron interactions via Bertini cascade model up to 10 GeV/c, a low energy parametrized model between beyond 10 and a quark gluon string model beyond 25 GeV/c. MAUS provides methods to set up the GEANT4 physical processes allows the user to control processes with datacard settings. Routines are also provided to the interface the internal geometry representation in MAUS with GEANT4 descriptions. Finally, MAUS provides routines to extract particle data from the GEANT tracks independent of the GEANT geometry with virtual planes.

3.3 Geometry

MAUS uses an on-line Configuration Database to store all of its geometries. These geometries have been extracted from CAD drawings which are based on the latest surveys and technical drawings available. The CAD drawings are translated to a geometry specific subset of XML, the Geometry Description Markup Language (GDML) prior to being recorded in the configuration database through the use of the Fast-RAD commercial software package.

The GDML formatted description contains the beam-line elements and the positions of the detector survey points. Beam-line elements are described using Tessellated solids to define the shapes of the physical volumes. The detectors themselves are described using an independently generated set of GDML files using GEANT4 standard volumes. An additional XML file is appended to the geometry description that assigns magnetic fields and associates the detectors to their locations in the GDML files. This file is initially written by the geometry maintainers and formatted to contain run specific information during download.

The GDML format has a number of benefits. The files can be read via a number of already existing libraries in GEANT4 and ROOT for the purpose of independent verification and validation. Because it is a subset of XML, the data contained in the GDML files are readily accessible through the application of the “libxml2” python extension. The GDML are in turn translated into the MAUS readable geometry files either by directly accessing the data using the python extension (which is the method applied to the detector objects) or through the use of EXtensible Stylesheet Language Transformations (XSLT).

3.4 Tracking, Field Maps and Beam Optics

MAUS tracking is performed by GEANT4 [11]. By default, MAUS uses 4th order Runge-Kutta for tracking, although other routines are available. 4th order Runge-Kutta has been shown to have very good precision relative to the MICE detector resolutions, even for step sizes of several cm.

Magnetic field maps are implemented as a series of overlapping regions, each of which contains a field. On each tracking step, MAUS iterates over the list of fields, transforms to the local coordinate system of the field map, and calculates the field. The field values are transformed back into the global coordinate system, summed and passed to GEANT4.

Numerous field types have been implemented within the MAUS framework. Solenoid fields can be calculated numerically from cylindrically symmetric 2D field maps, by taking derivatives of an on-axis solenoidal field or by using the sum of fields from a set of cylindrical current sheets. Pillbox fields can be calculated by using the Bessel functions appropriate for a TM010 cavity or by reading a cylindrically symmetric field map. Multipole fields can be calculated from a 3D field map, or by taking derivatives from the usual multipole expansion formulae. Linear, quadratic and cubic interpolation routines have been implemented for field maps.

Matrix transport routines for propagating particles and beams through these field maps have been implemented. Transport matrices are calculated by taking the numerical derivative of tracking output. These can be used to transport beam ellipses and single particles, for example enabling optics work, beam matching and so forth. Higher order transport routines are also available.

The accelerator modelling routines in MAUS have been validated against ICOOL and G4Beamline. The routines have been used to model a number of beamlines and rings, including the Neutrino

Factory front end.

3.5 Detector response and digitization

The modelling of the detector response and electronics enables MAUS to provide test data for reconstruction algorithms and estimate the errors introduced by the detector and its readout.

The interaction of particles in material is also performed by GEANT4. A “sensitive detector” class for each detector processes the hits in active volumes and stores relevant hit information such as the volume that was hit, the energy deposited and the time of the hit. The digitizers then simulate the detector response to these hits, modelling processes such as the photon yield of scintillator light, attenuation in light guides and the pulse shape in the electronics. The data structure of the outputs from the digitizers are designed to mock the output from the unpacking of the data from the DAQ.

4. Reconstruction

The reconstruction chain takes as its input either digitized hits from MC or DAQ digits from real data. Regardless, the detector reconstruction algorithms, by requirement and design, operate the same way on both MC and data. MAUS is currently capable of performing at least some reconstruction on all of the MICE detectors.

??? TODO: something about aims of the reconstruction – tracking, particle identification for purposes of physics goals [mention physics goals somewhere up top ?] ??? TODO: some intro/description of detectors with refs to pubs ??

4.1 Time of flight

There are three time-of-flight detectors in MICE and they serve to distinguish particles based on their times of flight. The detectors are made of plastic scintillator and in each station there are x and y planes with 6 or 10 slabs in each plane.

Each GEANT hit in the TOF is associated with a physical scintillator slab based on the geometry. The energy deposited by a hit in is first converted to units of photoelectrons. The photoelectron yield from a hit is attenuated by the distance from the hit to the PMT, then smeared by the photoelectron resolution. The yields from all hits in a given slab are then added and the summed photoelectron yield is converted to ADC counts.

The hit time is propagated to the PMTs at either end of the slab. The propagated time is then smeared by the PMT time resolution and converted to TDC counts. After converting the energy deposit to ADC and the time to TDC, the TDC values are “uncalibrated” so that at the reconstruction stage they can be corrected just as is done with real data.

The reconstruction proceeds in two main steps: 1) Slab-hit-reconstruction takes individual PMT digits and associates them to reconstruct the hit in the slab. If there are multiple hits associated with a PMT, the hit which is earliest in time is taken to be the real hit. Then, if both PMTs on a slab have hits, the slab is considered to have a valid hit. The TDC values are converted to time and the hit time and charge associated with the slab hit are taken to be the average of the two PMT times and charges respectively. In addition, the charge product of the PMT charges is also formed. 2) Finally, individual slab hits are used to form space-points. A space point in the TOF is a combination of x and y slab hits. All combinations of x and y slab hits in a given station are treated

as space point candidates. Calibration corrections, stored in the Configurations Database, are applied to these hit times and if the reconstructed space-point is consistent with the resolution of the detector, the combination is said to be a valid space point.

The TOF has been shown to provide good time resolutions, with errors at the 50 ps level. Improvements are being made to the calibration algorithms to improve the reconstructed resolutions and examine some suspected systematic effects.

??? TODO: Plots of resolution, time of flight??

4.2 Scintillating fiber trackers

The scintillating fiber trackers are the central piece of the reconstruction. There are two trackers, one upstream and the other downstream of an absorber, situated within solenoidal magnetic fields. The trackers are responsible for measuring the emittance before and after particles pass through the absorber.

Digits are the most basic unit fed into the main reconstruction module, each digit representing a signal from one tracker channel. Digits from adjacent channels are assumed to come from the same particle and are grouped to form clusters. Clusters from channels which intersect each other, in at least two planes from the same station, are used to form space-points, giving x and y positions where a particle intersected a station. Once space-points have been found, they are associated with individual tracks through pattern recognition (PR) (described in section 4), giving straight or helical PR tracks. These tracks, and the space-points associated with them, are then sent to the final track fit. To avoid biases that may come from space-point reconstruction, the Kalman filter uses only reconstructed clusters as input. [ACTION: Update/cite to tracker paper - AD]

??? TODO: Ref to Tracker paper ??? TODO: Plots of spatial, momentum resolutions??

4.3 KL calorimeter

Hit-level reconstruction of the KL is implemented in MAUS. Individual PMT hits are unpacked from the DAQ or simulated from MC and the reconstruction associates them to identify the slabs that were hit and calculates the charge and charge-product corresponding to each slab hit. [ACTION: Update - MB/JN]

??? TODO: Ref to Tracker paper ??? TODO: What plots??

4.4 Electron-muon ranger

Hit-level reconstruction of the EMR is now implemented in MAUS; the integrated ADC and time over threshold are calculated for each bar that was hit. [ACTION: Update - RA/FD]

??? TODO: Ref to EMR paper ??? TODO: Plots??

4.5 Cherenkov

The Ckov reconstruction takes the raw flash-ADC data, subtracts pedestals, calculates the charge and applies calibrations to determine the photoelectron yield. [ACTION: Update - DR/LC]

??? TODO: expand ??? TODO: plots??

4.6 Global reconstruction

The aim of the Global Reconstruction is to take the reconstructed outputs from individual detectors and tie them together to form a global track including a likelihood for various particle hypotheses.

Global track matching is performed by collating particle hits (TOFs 0, 1 and 2, KL and Ckov) and tracks (Trackers and EMR) from each detector using their individual reconstruction and combining them using a 4th order Runge-Kutta (RK4) method to propagate the particles' between these detectors.

Particle identification in MICE typically requires the combination of several detectors. Principally the time-of-flight between TOF detectors can be used to calculate velocity, which is compared with the momentum measured in the Trackers to calculate particle mass and hence particle type. For all but very low pT events, charge can be determined from the direction of helical motion in the Trackers. Additional information can be gleaned from the Ckov, KL and EMR detectors. The global particle identification framework is designed to tie this disparate information into a set of hypotheses of particle types, with an estimate of the likelihood of each hypothesis.

Global Reconstruction is performed outwards from the cooling channel; upstream of Tracker 1 (tracker_0 in MAUS) through TOF0; and downstream of Tracker 2 (tracker_1 in MAUS) through EMR. It is also available as a commissioning tool providing through-going tracks from TOF1 to EMR, in the absence of magnetic fields. During later Step IV analysis, time of flight between TOF1 and TOF2 will be included to improve the performance of the global reconstruction.

Efficiency and purity are returned for global track matching and PID along with global reconstruction efficiency (combined efficiency of the global reconstruction not including individual detector recon efficiencies) and global efficiency (combined efficiency of the global reconstruction, including the result of detector recon efficiencies) values.

4.6.1 Global Track Matching

Track points are matched to form tracks using a RK4 method. Initially this is done independently for the upstream and downstream (i.e. either side of the absorber) sections of the beamline. As the Trackers provide the most accurate position reconstruction, they are used as starting points for track matching, propagating hits outwards into the other detectors and then comparing the propagated position to the measured hit in the respective detector. The acceptance criteria for a hit belonging to a track is an agreement within the fundamental limit (i.e. half the detector granularity) plus an allowance for multiple scattering. Track matching is currently performed for all TOFs, KL and EMR.

The RK4 propagation requires the mass and charge of the particle to be known, therefore, it is necessary to perform track matching for all realistically possible particle types (muons, pions, and electrons), though typically the Tracker reconstruction will provide a charge hypothesis to exclude a charge-sign beforehand. Tracks for all possible PID hypotheses are then passed to the PID algorithms.

Efficiency and purity of the track matching are evaluated on a per-detector basis by comparing MC tracks that could have been matched—which requires MC hits existing in both the respective tracker (upstream or downstream) and the detector for which the efficiency is being calculated. In order for an event to be used in this evaluation, the tracker hits have to have been reconstructed

correctly or in such a way so as to not alter the path of the particle in the tracker, as otherwise track matching can't reliably be performed. Checks are then performed to determine whether the correct hit has been matched.

For use during the alignment run, track matching between upstream and downstream tracks in a no-absorber no-field scenario has been implemented using a cut on the TOF1-TOF2 time-of-flight as the principal matching criteria ($z/c < dt < z'/\sim 0.6c$). Track matching will be extended in the future to be able to handle triggers containing more than one particle.

Track fitting will be performed on tracks returned from the PID in the future.

??? TODO: Plots??

4.6.2 Global PID

The Global PID in MAUS uses a log-likelihood method to determine the particle identification (PID) of a global track. It is based upon a framework of PID variables; quantities that can be used to distinguish between different particle species. These variables are used with Monte Carlo tracks to produce probability density functions (PDFs) of their values. PID variables for a real data track are then compared to the corresponding PDFs, in order to obtain a set of likelihoods for the potential PIDs of the track.

Track reconstruction will pass to the PID a number of potential tracks, each reconstructed for a given particle hypothesis. The PID will take each of these tracks in turn, and determine the most likely PID of that track. This process follows a series of steps:

1. Each track is copied into an intermediate track, and the PID of the copy is set to zero (undefined).
2. For each potential PID hypothesis x , the log-likelihood LL_x is calculated using the PID variables.
3. The track is assigned an object containing the log-likelihood for each PID hypothesis (for reference).
4. From the log-likelihoods, the confidence level for a track having a PID x (CL_x) is calculated by $CL_x = (LL_x)/(\sigma_i LL_i)$
5. The PID of the track is set to the particle hypothesis that gave the best CL.

If the PID determined above matches the PID that the track reconstruction had assigned to the track, then that track (with its PID) is considered the correct track, to then be passed back to track reconstruction for final fitting. Otherwise, the PID is unable to assign a PID to the particle and does not return a final track.

PID will be used for both global alignment during commissioning, and running of Step IV. However, the different requirements of these phases of the experiment necessitate two different sets of PID variables. During Step IV, the requirement for independence between upstream and downstream track reconstruction means that the PID variables used are also separated into two sets, one using the upstream detectors, the other the downstream detectors. The variables included in MAUS are:

PID class	Variable name	Definition
PIDVarA	diffTOF1TOF0	Uses the upstream time of flight, between TOF0 and TOF1. This variable is beam dependent and so is best used during offline data analysis where PDFs can be produced for specific beam settings.
PIDVarB	diffTOF0TOF1vsTrackerMom	Uses upstream time of flight, and momentum as measured in the upstream Tracker.
PIDVarC	KLChargeProdvsDSTrackerMom	Uses the KL ADC charge product and the momentum measured in the downstream Tracker.

The following variables are under development:

PID class	Variable name	Definition
PIDVarD	KLADCCChargeProduct	Uses the KL ADC charge product. This variable is beam dependent, and so is best used during offline data analysis.
PIDVarE	EMRRangevsDSTrackerMom	Uses the range of the particle as measured in the EMR, and the momentum measured in the downstream Tracker.
PIDVarF	EMRPlaneDensityvsDSTrackerMom	Uses the plane hit density in the EMR, and the momentum measured in the downstream Tracker.

For global alignment and commissioning, where there will be no magnetic field, the Trackers will not be able to provide momentum information, and TOF0 cannot be used (as a series of quadrupole magnets sit between TOF0 and TOF1 which will cause an unknown perturbation to the field when not running). However, alignment will use tracks that pass through the length of the cooling channel, and so PID (variables that incorporate both upstream and downstream detectors can be used). The PID variables for use during commissioning are:

PID class	Variable name	Definition
ComPIDVarA	diffTOF2TOF1	Uses the time of flight between TOF1 and TOF2. Beam dependent.
ComPIDVarB	KLChargeProdvsDiffTOF1TOF2	Uses the KL ADC charge product, and the time of flight between TOF1 and TOF2.
ComPIDVarC	ComKLADCChargeProduct	Uses the KL ADC charge product . Beam dependent.

the following variables are under development:

PID class	Variable name	Definition
PIDVarD	EMRRangevsDiffTOF1TOF2	Uses the range of the particle as measured in the EMR, and the time of flight between TOF1 and TOF2.
PIDVarE	EMRPlaneDensityvsDiffTOF1TOF2	Uses the plane hit density in the EMR, and the time of flight between TOF1 and TOF2.

Efficiency and purity studies are used to improve the performance of the PID. This is performed on both a variable-by-variable basis, and globally across the variables. The efficiency of the PID is defined as $eff = \text{correct PID tracks} / (\text{MC tracks} - \text{unsuitable tracks})$ whereby an unsuitable track is one that does not fulfill the criteria required to perform PID on it (i.e. missing detector information). The purity of the PID is defined as $p = \text{correct PID tracks} / \text{tracks assigned a PID}$.

4.7 Online reconstruction

During data taking, it is essential to visualize a detector's performance and have diagnostic tools to identify and debug unexpected behavior. This is accomplished through summary histograms of high and low-level reconstructions from detectors. These are available for the TOF, Cherenkov, KL, EMR and Trackers.

For online reconstruction, MAUS uses a distributed processing model to enable a scalable reconstruction of the MICE dataset. Raw data is passed to a networked message queue for multiprocessing across multiple CPUs and servers. Reconstructed data is handed to another message queue. Histogramming routines pick data from this second message queue and collate it into histograms, which are written to disk. A web-based visualisation tool enables viewing of the histograms. A production version of the online reconstruction is available. Further development work is underway on the user interface and some of the backend infrastructure. Though the framework for the online reconstruction is based on distributed processing of spills, the reconstruction modules are the same as those used for offline processing.

An event display summarising global reconstruction data will be implemented when the global reconstruction is complete. This will enable visualisation of the phase space distribution of the

beam at various points along the beamline, together with a comparison of the nominal beam envelope propagation. The event display is intended to enable online validation of the behaviour of accelerator components, by comparing propagation of the beam envelope with the detected beam parameters.

??? TODO: Redo re new Yordan framework

5. Summary

??? TODO:

Acknowledgments

Aacknowledgments.

References

- [1] G. Gregoire et al., *Proposal to the Rutherford Appleton Laboratory: An International Muon Ionization Cooling Experiment (MICE)*, **Tech. rep. (2003)** <http://mice.iit.edu/micenotes/public/pdf/MICE0021/MICE0021.pdf>
- [2] Geer S 2009 “Muon Colliders and Neutrino Factories,” *Ann. Rev. Nucl. Part. Sci.* **59** 345–367.
- [3] S. Choubey et al., *International Design Study for the Neutrino Factory, Interim Design Report, FERMILAB-PUB-11-581-APC, 2011*
- [4] <http://launchpad.net>
- [5] <http://bazaar.canonical.com/>
- [6] <http://www.gridpp.ac.uk/>
- [7] C.D. Tunnell and C.T. Rogers, *MAUS: MICE Analysis User Software, IPAC (2011)*;
<http://micewww.pp.rl.ac.uk/projects/maus>
- [8] <http://jenkins-ci.org>
- [9] <http://g4beamline.muonsinc.com>
- [10] R.C. Fernow, *ICOOL: A simulation code for ionization cooling of muon beams, Proc. 1999 Particle Accelerator Conference, New York (1999)*.
- [11] S. Agostinelli, et al., *GEANT4 - A Simulation Toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250-303
- [12] <http://json.org>
- [13] R. Brun and F. Rademakers, *ROOT - An Object Oriented Data Analysis Framework*, *Nucl. Instr. Meth. A* **389**, 81-86 (1997)




Figure 3. The MAUS output structure for a spill event. The top label in each box is the name of the C++ class and the bottom label is the json branch name. If a [] is shown, this indicates that child objects are array items.

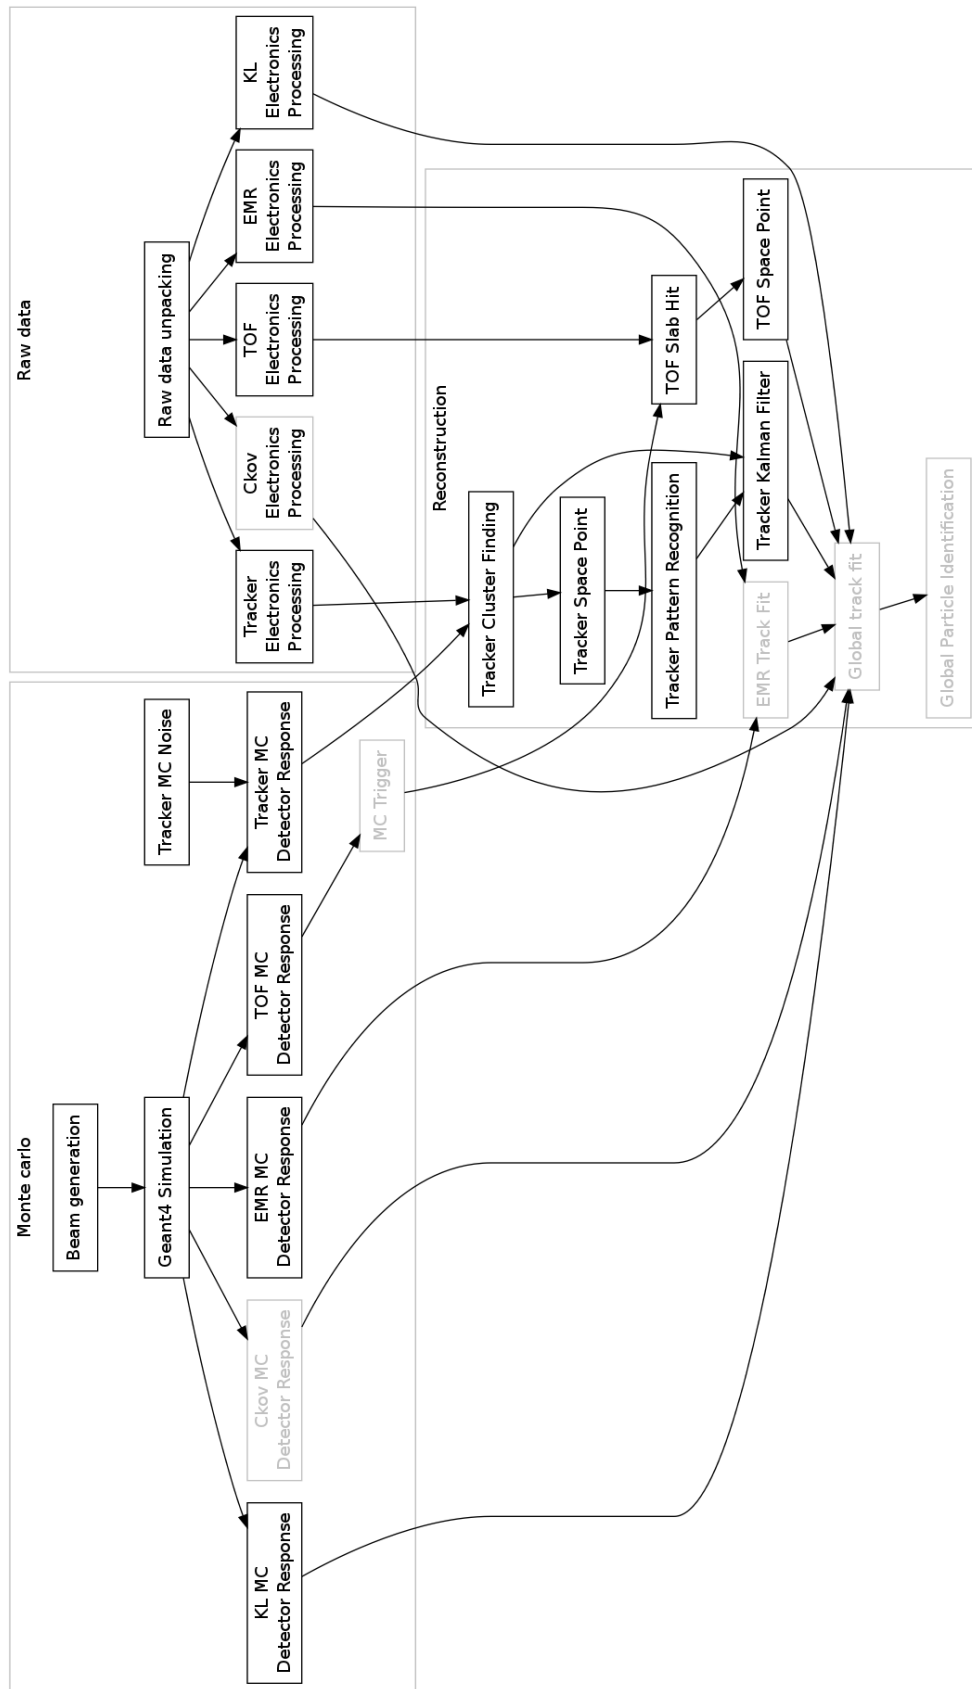


Figure 4. Data flow for the MAUS project. Items shown in black have some production implementation, although not necessarily the final one. Those shown in grey are under development.